Professional Certificate in AI for Chemical Process Engineering

Fundamentals of Machine Learning

Fundamentals of Machine Learning Glossary

Activation Function: A mathematical function that introduces non-linearity to neural networks, allowing them to learn complex patterns in data. Common examples include sigmoid, tanh, ReLU, and softmax functions.

Artificial Intelligence (AI): The simulation of human intelligence processes by machines, especially computer systems. AI encompasses tasks such as learning, reasoning, problem-solving, perception, and language understanding.

Backpropagation: An algorithm used in training artificial neural networks to update the weights of the network by calculating the gradient of the loss function with respect to the weights.

Batch Size: The number of training examples utilized in one iteration (epoch) of training in a machine learning model.

Bias: A term in machine learning that represents a model's tendency to consistently learn the wrong thing by not taking into account all the information in the data.

Classification: A type of supervised learning where the goal is to predict the categorical class labels of new instances based on past observations.

Clustering: An unsupervised learning technique used to group data points into clusters based on similarities in the data.

Convolutional Neural Network (CNN): A type of deep neural network commonly used for analyzing visual imagery. CNNs have convolutional layers that automatically learn hierarchical patterns in the data.

Cross-Validation: A technique used to evaluate the performance of a machine learning model by splitting the data into multiple subsets for training and testing.

Deep Learning: A subfield of machine learning that focuses on learning representations of data through multiple layers of neural networks.

Dimensionality Reduction: The process of reducing the number of input variables in a dataset by obtaining a set of principal variables.

Dropout: A regularization technique used in neural networks to prevent overfitting by randomly disabling a

fraction of neurons during training.

Ensemble Learning: A machine learning technique that combines multiple models to improve the overall performance and robustness of the system.

Feature Engineering: The process of selecting, combining, and transforming variables in a dataset to improve the performance of machine learning algorithms.

Gradient Descent: An optimization algorithm used to minimize the loss function by iteratively moving in the direction of the steepest descent.

Hyperparameter: Configurable variables that dictate the learning process of a machine learning algorithm, such as learning rate, batch size, and number of epochs.

Kernel: A function used to transform input data into a higher-dimensional space to make it linearly separable for classification tasks.

Logistic Regression: A statistical model used for binary classification that predicts the probability of a binary outcome.

Loss Function: A function that quantifies the difference between the predicted values of a model and the actual values in the training data.

Neural Network: A computational model inspired by the human brain that consists of interconnected nodes (neurons) organized in layers.

Overfitting: A phenomenon in machine learning where a model learns noise in the training data instead of the underlying patterns, leading to poor generalization on unseen data.

Principal Component Analysis (PCA): A dimensionality reduction technique that transforms the data into a new coordinate system to capture the maximum variance in the data.

Reinforcement Learning: A type of machine learning where an agent learns to make decisions by interacting with an environment and receiving rewards or penalties.

Regularization: Techniques used to prevent overfitting by adding a penalty term to the loss function that discourages complex models.

Regression: A type of supervised learning where the goal is to predict continuous output values based on input features.

Stochastic Gradient Descent: A variant of the gradient descent algorithm that updates the weights of a model using a small subset of training examples at a time.

Support Vector Machine (SVM): A supervised learning algorithm used for classification and regression tasks by finding the hyperplane that best separates the classes in the data.

Underfitting: A situation where a machine learning model is too simple to capture the underlying patterns in the data, leading to poor performance on both training and test sets.

Unsupervised Learning: A type of machine learning where the goal is to find hidden patterns or intrinsic structures in the data without using labeled examples.

Validation Set: A subset of data used to tune hyperparameters and evaluate the performance of a machine learning model during training.

Variance: A measure of how much the predictions of a model vary for different training sets, indicating the model's sensitivity to changes in the training data.

Word Embedding: A technique used to represent words as dense vectors in a continuous vector space, capturing semantic relationships between words based on their context.

This glossary provides a comprehensive overview of key concepts and terms in machine learning essential for the Professional Certificate in AI for Chemical Process Engineering. By understanding these fundamental principles, learners can effectively apply machine learning techniques to solve real-world problems in their field.